



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

THESIS

**LOWER TOTAL COST OF OWNERSHIP OF ONE-NET BY
USING THIN-CLIENT DESKTOP DEPLOYMENT AND
VIRTUALIZATION-BASED SERVER TECHNOLOGY**

by

Mytiec Lam

September 2010

Thesis Co-Advisors:

Thomas V. Huynh
John S. Osmundson

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE September 2010	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE Lower Total Cost of Ownership of ONE-NET by Using Thin-Client Desktop Deployment and Virtualization-Based Server Technology		5. FUNDING NUMBERS	
6. AUTHOR(S) Mytie Lam			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A		10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol Number _____ N/A			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited		12b. DISTRIBUTION CODE A	
13. ABSTRACT (maximum 200 words) The U.S. Navy (USN) has a large enterprise network outside the continental U.S. (OCONUS) that is complex and expensive to maintain. The problem addressed by this thesis is to determine which alternative desktop deployment technology is more cost effective over time while maintaining the users' operational requirements. The USN is conducting a technology refresh of its OCONUS navy enterprise network (ONE-NET) with thick-client desktop computers. This thesis proposes an alternative solution using thin-client desktops with data center server virtualization-based technology as a lower cost option. To back up this claim of lower cost, an analysis was carried out to determine the total ownership costs (TCO) of both the current thick-client and proposed thin-client solutions. A cost per seat (CPS) model developed by Naval Network Warfare Command (NNWC) was used to calculate major cost components—labor, hardware, software, and transport, while a VMware tool was used to calculate power and cooling costs for both solutions. In addition, VMware provided a cost estimate for the upfront hardware and software licensing costs needed to support the virtualization support for the thin-clients solution. The conclusion of the TCO comparison is that, for the 27,284 users, the thin-client solution would save \$238 million over seven years.			
14. SUBJECT TERMS Thick-Client, Thin-Client, Server Virtualization, Total Cost of Ownership, OCONUS Navy Enterprise Network, ONE-NET			15. NUMBER OF PAGES 79
16. PRICE CODE			
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

LOWER TOTAL COST OF OWNERSHIP OF ONE-NET BY USING THIN-CLIENT DESKTOP DEPLOYMENT AND VIRTUALIZATION-BASED SERVER TECHNOLOGY

Mytiec Lam
Civilian, Department of the Navy
B.S., University of California at Davis, 1985

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN SYSTEMS ENGINEERING MANAGEMENT

from the

NAVAL POSTGRADUATE SCHOOL
September 2010

Author: Mytiec Lam

Approved by: Thomas V. Huynh, PhD
Thesis Co-Advisor

John S. Osmundson, PhD
Thesis Co-Advisor

Cliff Whitcomb, PhD
Chairman, Department of Systems Engineering

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The U.S. Navy (USN) has a large enterprise network outside the continental U.S. (OCONUS) that is complex and expensive to maintain. The problem addressed by this thesis is to determine which alternative desktop deployment technology is more cost effective over time while maintaining the users' operational requirements. The USN is conducting a technology refresh of its OCONUS navy enterprise network (ONE-NET) with thick-client desktop computers. This thesis proposes an alternative solution using thin-client desktops with data center server virtualization-based technology as a lower cost option. To back up this claim of lower cost, an analysis was carried out to determine the total ownership costs (TCO) of both the current thick-client and proposed thin-client solutions. A cost per seat (CPS) model developed by Naval Network Warfare Command (NNWC) was used to calculate major cost components—labor, hardware, software, and transport, while a VMware tool was used to calculate power and cooling costs for both solutions. In addition, VMware provided a cost estimate for the upfront hardware and software licensing costs needed to support the virtualization support for the thin-clients solution. The conclusion of the TCO comparison is that, for the 27,284 users, the thin-client solution would save \$238 million over seven years.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	BACKGROUND	1
B.	OBJECTIVE	4
C.	RESEARCH QUESTIONS	4
D.	RESEARCH BENEFIT.....	4
E.	SCOPE AND METHODOLOGY	5
1.	Scope.....	5
2.	Methodology	5
F.	THESIS ORGANIZATION.....	5
II.	CENTRALIZED SERVERS AND THIN-CLIENT COMPUTING.....	7
A.	DEFINITIONS AND THEORY	7
B.	THIN-CLIENT AND SERVER TECHNOLOGY EVOLUTION	8
1.	First Generation.....	8
2.	Second Generation	10
3.	Third Generation	11
C.	CHAPTER SUMMARY	11
III.	KEY ARCHITECTURE EVALUATION FACTORS	13
A.	OVERVIEW	13
B.	INFRASTRUCTURE AND NETWORK SERVICES	13
1.	Footprint and Facilities	14
2.	Networking Resources	18
C.	SOFTWARE DELIVERY AND MANAGEMENT.....	21
1.	Baseline Configuration Management.....	21
2.	Image Build Process and Administration	22
3.	Application Management and Licensing	24
D.	ENTERPRISE DEPLOYMENT AND INFORMATION ASSURANCE	27
E.	SERVICE LEVEL PERFORMANCE.....	33
F.	CHAPTER SUMMARY	33
IV.	TOTAL COST OF OWNERSHIP	35
A.	TOTAL COST OF OWNERSHIP CONCEPT.....	35
B.	MAJOR COST COMPONENTS AND ASSUMPTIONS.....	36
1.	Labor Costs.....	36
2.	Hardware Costs.....	36
3.	Software Costs	36
4.	Transport Costs.....	37
C.	COST ANALYSIS AND TCO COMPARISON	37
1.	Thick-Client Cost Analysis.....	38
2.	Thin-Client Cost Analysis	41
3.	TCO Comparison.....	44
D.	CHAPTER SUMMARY	45

V. CONCLUSION	49
APPENDIX A. CPS MODEL CALCULATION RESULTS FOR FY08 UNADJUSTED THICK-CLIENT SOLUTION.....	51
APPENDIX B. REGIONAL ROLLUP CALCULATIONS FOR TCO OF THICK- CLIENT SOLUTION	53
APPENDIX C. VMWARE CAPITAL AND OPERATING INVESTMENT COSTS FOR VIRTUALIZATION	55
LIST OF REFERENCES	57
INITIAL DISTRIBUTION LIST	59

LIST OF FIGURES

Figure 1.	ONE-NET Sites (From SSC PAC, 2010a)	2
Figure 2.	TNOSC and LNSC Unclassified Server Farm Layout (From CSC, 2009)	15
Figure 3.	High-Level TNOSC Data Center Infrastructures (From CSC, 2009).....	17
Figure 4.	VMware® Building Block for 1000 Virtual Desktop Users (From VMware, 2010).....	18
Figure 5.	Networking Server Farm with Six VLANs (From CSC, 2009).....	20
Figure 6.	WBSC118 Image Build Architecture (From Lelfwitch, 2009)	23
Figure 7.	Optional WBSC Baseline Image Build Progressions (From Lelfwitch, 2009)	24
Figure 8.	Altiris® Enterprise Deployment Architecture (From Martinez, 2010)	28
Figure 9.	Staging VLAN Architecture (From Hanada, 2009).....	29
Figure 10.	Post Imaging PC Patching Process (From Graham, 2009)	30
Figure 11.	ONE-NET Windows Patch Management Process (From SSC PAC, 2009b)..	32

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	Required Software Components to Create WBSC Image (From Lelfwitch, 2009)	22
Table 2.	Baseline Applications on WBSC Image Build (SSC PAC, 2009a).....	25
Table 3.	Original Inflation Index (From NNWC, 2005).....	38
Table 4.	Four Major Cost Components for Thick-Client Solution	39
Table 5.	FY08 Regional Rollup for Thick-Client Solution (Appendix B)	39
Table 6.	Annual Cost of Operating and Cooling Power on Thick-Client Solution (From VMware, 2010)	41
Table 7.	Four Major Cost Components for Thin-Client Solution.....	42
Table 8.	Annual Cost of Power and Cooling on Thin-Client Solution (From VMware, 2010).....	43
Table 9.	IT Capital and Operating Investment (From VMware, 2010)	44
Table 10.	TCO Comparison between Between Thick-Client and Thin-Client Solutions	45
Table 11.	FY 08 Unadjusted Thick-Client Solution Cost Results	51
Table 12.	FY 05 Through FY 11 Regional Rollup Costs	53
Table 13.	VMware Virtualization Capital and Operating Investment Costs	55

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

ARPA	Advanced Research Projects Agency
BAN	Base Area Networks
BLII	Base Level Information Infrastructure
BSD	Berkeley Software Distribution
CANES	Consolidated Afloat Networks and Enterprise Services
CIO	Chief Information Officer
CONOPS	Concepts of Operations
CONUS	Continental United States
CPS	Citrix Presentation Server
CSRG	Computer Systems Research Group
DISN	Defense Information System Network
DON	Department of Navy
EDS	Electronic Data Systems
GUI	Graphical User Interface
IP	Internet Protocol
LNSC	Local Network Support Center
LOB	Line-of-Business
MAN	Metropolitan Area Networks
MTS	Microsoft Terminal Services
NNWC	Naval Network Warfare Command
NGEN	Next Generation Enterprise Network
NMCI	Navy Marine Corps Intranet
NNE	Naval Network Environment
NSC	Network Support Center
OCONUS	Outside Continental United States
ONE-NET	OCONUS Navy Enterprise Network
OS	Operating Systems
PC	Personnel Computer

PEO-EIS	Program Executive Office – Enterprise Information Systems
POM	Program Objective Memorandum
RAM	Random Access Memory
TCO	Total Cost of Ownership
TNOSC	Theater Network Operations Security Center
VDM	Virtual Desktop Machine
WAN	Wide Area Networks
XDMCP	X-Windows Display Manager Control Protocol

EXECUTIVE SUMMARY

The U.S. Navy (USN) has a large enterprise network outside the continental U.S. (OCONUS), which is complex and expensive to maintain. The OCONUS navy enterprise network (ONE-NET) covers three major regions identified as the island nation of Bahrain, Far East, and Europe. There are nine sites in the Far East, four sites in Europe, and one site in Bahrain that comprise the network support centers (NSC) of the ONE-NET enterprise network. Starting in 2010 through 2016, the USN will be refreshing the ONE-NET technology, after which it will be replaced by a new next generation enterprise network (NGEN). In this technology refresh, the USN is in the process of deploying thick-client desktop computers to 27,284 users. However, thick-client technology is not the only available solution as there are several thin-client solutions as well that can potentially meet the operational requirements of ONE-NET at a lower cost. Historically, thin-client technology has gone through several generations of improvement in capabilities and has recently shown that it can provide the same user experience as the thick-client solution. Therefore, it raises the question of whether the thick-client solution is the most cost effective one. Given the obvious goal of any organization to save money where it can on expenses, the challenge is to properly plan and execute a long-term service capability such as ONE-NET, which will have the lowest total ownership costs (TCO).

The problem addressed by this thesis is: which alternative desktop technology is more cost effective over the next seven years (until NGEN replaces ONE-NET), while maintaining the users' operational requirements? This thesis proposes that an alternative solution using thin-client desktops with data center server virtualization-based technology could be the lower cost option over the current thick-client solution. To support this claim of lower TCO, a study was done on the technology and operation of the current ONE-NET in order to get an understanding of the major cost contributing factors, and an analysis was then done to obtain the TCO for both thick-client and thin-client solutions. A cost per seat (CPS) model developed by Naval Network Warfare Command (NNWC) was used to calculate the major cost components for labor, hardware, software, and

transport using various pricing and labor rates as inputs to the model. Both solutions assume 27,284 user seats and a period of seven years. Since the original model was done for a period starting in 2005, an inflation index was used to calculate the inflation offset of 10.7% for 2010, and it was applied to the model in order to shift the calculations to cover the period from 2010 through 2016. The results of the CPS model show that the labor and hardware costs for the thick-client are, respectively, \$545 million and \$97 million. While the labor and hardware costs for the thin-client solution are, respectively, \$335 million and \$39 million. In addition, VMware provided a cost estimate for the upfront hardware and software licensing costs needed to support the virtualization support for the thin-clients solution, which is not a cost factor for the thick-client solution. This upfront cost is \$41 million over the five-year period of rolling out the technology refresh, which is the same period for the thick-client solution. Having the same schedule for comparison also provides a confirmation that the proper pricing is being used as it is normally tied to a schedule—the shorter time period often raises the cost for the same fixed job. Finally, the power and cooling costs, obtained with a VMware tool, for both the thick-client and thin-client solutions indicate that the thin-client solution saved \$11.6 million.

The conclusion of the TCO comparison is that for the 27,284 users, the thin-client solution would save the sponsoring organization \$238 million over seven years. This is a significant amount considering that an acquisition program of over \$100 million is considered a major acquisition program, and this is just the savings amount. Therefore, the results of this thesis could potentially benefit the U.S. Navy's program executive office—enterprise information systems (PEO-EIS), which is the ONE-NET acquisition program sponsor.

Further research is recommended in the latest technology development area of cloud computing, offering an interesting topic for future work on how the concept relates to enterprise networks and TCO. The area of information assurance (IA), a challenging problem for cloud computing, is also a potential topic for future research.

ACKNOWLEDGMENTS

I am very grateful for the help and advice Dr. Tom Huynh and Dr. John Osmundson have provided, and I appreciate their time spent in helping me complete this goal of mine. If it were not for their support and encouragement, I would not be able to fulfill my NPS program requirements.

Also, I would like to thank Mr. Jack Cabana, SPAWAR Systems Center Pacific Code 40 Department Head, for providing me with the opportunity to pursue my goal of attaining an MS in Systems Engineering Management at NPS.

I am thankful to Mr. Michael Ferguson, SPAWAR Systems Center Pacific Code 712 Chief Engineer, for his technical proof reading and encouragement.

Finally, I thank my family for their support and encouragement over the past two years.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

The purpose of this chapter is to provide a brief description of the background that motivated this thesis, including the introduction of the subject matter and some terminology used to describe the technology of thin-client computing. A short description of the factors and benefits derived from the proposed use of a thin-client solution is provided that was used for calculating the Total Cost of Ownership (TCO). In addition, the thesis objective is defined, followed by the research questions, research benefits, and an overview of the scope and methodology used in this thesis.

A. BACKGROUND

PC Magazine defines an Enterprise Network as “[a] geographically dispersed network under the jurisdiction of one organization. It often includes several different types of networks and computer systems from different vendors” (*PC Magazine*, 2010).

In 2000, the Department of Navy (DON) awarded an outsourcing contract to Electronic Data Systems (EDS) for a consolidated DON enterprise network called the Navy Marine Corps Intranet (NMCI). The NMCI contract was set up to provide the majority of Information Technology (IT) services for the DON, including the United States Navy and Marine Corps. However, NMCI was limited to the Continental United State (CONUS), so another contract was awarded in 2001 to General Dynamics, Government Systems Corp., for the Base Level Information Infrastructure (BLII) improvements at the Navy’s Outside Continental United States (OCONUS) installations. Unlike NMCI, BLII is owned by Program Executive Office—Enterprise Information Systems (PEO-EIS) and operated by the Navy Network Warfare Command (NNWC). The consolidated DON enterprise network at the OCONUS was called as the OCONUS NMCI in 2001 and renamed the OCONUS Navy Enterprise Network in 2002, which is abbreviated as ONE-NET. The term ONE-NET will be used from here on.

From a global perspective, Figure 1 shows ONE-NET’s distribution coverage of the OCONUS naval bases, posts, camps, and stations, with their geographic locations around the world. Nine sites in the Far East, four sites in Europe, and one site in Bahrain

comprise the Network Support Centers (NSC) of the ONE-NET enterprise network. The three central sites are known as the ONE-NET Theater Network Operations Security Centers (TNOSCs), which are located in Yokosuka, Japan; Naples, Italy; and the island nation of Bahrain. The remaining fourteen are known as Local Network Support Centers (LNSCs). All of the ONE-NET locations are logically connected to each other via the Defense Information System Network (DISN) and have varying ranges of bandwidth. As its connectivity is not provided, DISN is referred to as a ‘Cloud’, an industry term used to describe a logical network. Throughout the regions covered by DISN, there are Base Area Networks (BAN), Metropolitan Area Networks (MAN) and Wide Area Networks (WAN).




Figure 1. ONE-NET Sites (From SSC PAC, 2010a)

The DON Chief Information Officer (CIO) is responsible for developing the vision, strategy, and Concept of Operations (CONOPS) for DON's future Naval Network Environment (NNE). The NNE is composed of four major network components: Next Generation Enterprise Network (NGEN) as the follow-on to NMCI, BLII/ONE-NET,

Consolidated Afloat Networks and Enterprise Services (CANES) as the shipboard component, and the other remaining legacy networks.

The motivation for this thesis is to explore the financial opportunity of leveraging the current ONE-NET enterprise networking technology refresh by replacing the thick-client Personal Computers (PC) with thin-client devices. A thick-client or fat-client is a full-featured PC that is connected to a network, while thin-client lacks hard drives and other features.

Thin-client computing has been around since the beginning of computers, but it has gone through some cycles where hardware and software do not always match in their capabilities. Most often, the software is more advanced than the supporting hardware. Starting with mainframe computers and directly attached terminals, the processing then shifted to the smaller and cheaper PC. As lower-cost workstations took over the computationally intensive portion of the mainframe computer processing, a new class of low-cost terminals evolved that provided access to remote centralized workstations. However, the terminals quickly lost their appeal because the early Internet could not support the bandwidth required to support the full Graphical User Interface (GUI) of the new terminals. The chosen solution between PCs and terminals swung back and forth a few more times before achieving today's satisfactory performance of centralized servers and thin-client computing technology.

This thesis is focused only on TCO of ONE-NET, because the other NNE network components, for which the technology is also applicable, are currently outsourced and their implementation is not under DON control. Although TCO is very complex and difficult to prove as an absolute cost that takes into account every factor that contributes to the total cost, this thesis focuses on the major cost contributing factors for the acquisition of the ONE-NET technology and services. The major cost contributing factors that were analyzed are labor, hardware, software, and transport.

TCO provides a long-term measure of cost, and this thesis covers a seven-year period to cover the support required for ONE-NET.

B. OBJECTIVE

The purpose of this thesis is to demonstrate how Total Cost of Ownership (TCO) of the Outside Continental U.S. (OCONUS) Navy Enterprise Network (ONE-NET) can be reduced by deploying thin-client devices, instead of full desktop PCs, and leveraging the centralized server virtualization technology being deployed in the data centers. Data centers are facilities used to house systems such as computer servers, associated telecommunication routers, storage components, etc.

C. RESEARCH QUESTIONS

The primary questions are:

1. What is the Total Cost of Ownership (TCO) for the current OCONUS Navy Enterprise Network (ONE-NET) over the next seven years? Note that the seven-year period enables the alignment of the thesis analysis with the planned lifespan of ONE-NET before it will be replaced with NGEN by NNWC.
2. Can the TCO be reduced by using a thin-client desktop deployment and virtualization-based server technology?

Answering the first question amounts to deriving the TCO of the current thick-client solution, which is a proven technology, whereas the proposed alternative thin-client solution adds some technical risk. The second question is responded to by answering the following detailed questions.

1. Will replacement of the current thick-client desktops with thin-client desktops meet the operational requirements of ONE-NET?
2. Will the TCO of ONE-NET over the next seven years be low enough to justify the additional risk and up-front investment in deploying the thin-client desktops?
3. Are there any other compelling advantages to deploying thin-client desktops and virtualization-based server technology over the current thick-client desktops?
4. Can the deployment of thin-client desktops be done during the current seven year technology support period and still lower the TCO?
5. What would be the recommended deployment methodology that will achieve the lower TCO?

D. RESEARCH BENEFIT

Enabling a reduction of the TCO, the use of thin-client desktops and centralized server virtualization technology, espoused in this thesis, will potentially benefit the U.S.

Navy's Program Executive Office–Enterprise Information Systems (PEO-EIS), the sponsor of ONE-NET.

E. SCOPE AND METHODOLOGY

1. Scope

The focus of this thesis is on the TCO comparison between the current thick-client solution and the proposed thin-client alternative, which also leverages the centralized server virtualization technology. The comparison attributes are technical capability, deployment schedule feasibility, and the TCO. The TCO accounts for costs of labor, hardware, software, and data transport as well as the additional cost of power and cooling.

2. Methodology

The methodology used in this research involves the following.

1. Conducting a review of client technology evolutions;
2. Identifying and analyzing key architecture evaluation factors that impact the cost and performance of implementing the alternative thin-client solution;
3. Identifying and justifying some assumptions required by the calculation model inputs;
4. Performing quantitative cost calculations using the major contributed factors for the TCO calculation model and comparison;
5. Formulating a recommendation regarding the implementation of the alternative solution that could benefit the ONE-NET acquisition program office.

F. THESIS ORGANIZATION

The thesis is organized into five chapters including this Chapter I, which provides the introduction. Chapter II discusses the topic of centralized servers and thin-client computing technology. The analysis and identification of the key architecture evaluation factors is done in Chapter III. Chapter IV presents and calculates the TCO of both alternative solutions for comparison. Finally, Chapter V provides the conclusions and recommendations based on the results of this research.

THIS PAGE INTENTIONALLY LEFT BLANK

II. CENTRALIZED SERVERS AND THIN-CLIENT COMPUTING

The purpose of this chapter is to provide the definitions and the theory of client-server computing model. A historical review on thin-client and server technology evolution is provided as a means to gain understanding of the types of issues and improvements that result from one technology generation to the next.

A. DEFINITIONS AND THEORY

A client-server model includes at a minimum two processes—one for the client and one for the server—and a communications channel between them. The client and server processes can reside on a multiprocessing computer or on two different computers. Communications among multiple computers on a network are effected by two major parts. The first part is the Berkeley Software Distribution (BSD) Sockets, which is “a UNIX operating system derivative developed and distributed by the Computer Systems Research Group (CSRG) of the University of California, Berkeley, from 1977 to 1995” (McKuick, 1999). The BSD Sockets are the mechanism by which computers interface with the network. The second part is the physical network with protocols that facilitate communications over the network, which was funded by the U.S. Advanced Research Projects Agency (ARPA) and later evolved into the Internet. Although there are numerous networking protocols, the predominant one today is the Internet Protocol (IP) along with many higher-level protocols that were developed by ARPA. In summary, the client and server use Sockets to communicate with each other over the network using IP-based protocols (Bach, 1986).

Now, what processing should the client do and what processing should the server do? The answer to this question is related to the capabilities and costs of the available technology. Client hardware, along with its supporting software, comes in several configurations that are based on the implemented client server model for distributing the processing loads between the client and the server. A thick-client or fat-client is a full-featured computer (PC) that is connected to a network. “Unlike thin-clients, which lack hard drives and other features, thick-clients are functional whether they are connected to

a network or not” (Tech Terms, 2010). Thin-client technology varies between the available hardware resources and the protocols they rely on for communicating with the servers. The expected performance also determines the resources needed by the servers and network infrastructure to handle the required communications bandwidth. Several variants of thin-client technologies include truly thin-client or zero-client, which is akin to the terminal concept where all the processing except the Graphical User Interface (GUI) is done by the server.

B. THIN-CLIENT AND SERVER TECHNOLOGY EVOLUTION

It is important to understand the thin-client technology evolution as it provides essential lessons learned that help in designing future enterprise networks. The following is a short overview of the general three informal generations of thin-client and corresponding server technologies. The term ‘generation’ here means a generalized grouping over a time period in which the particular thin-client and server technology is the predominant solution. Note that the term ‘thin-client’ post-dates the terminal days. In a sense, the so-called “dumb terminals” were actually an early version of thin clients, which were used to interface with mainframe computers by sending keystrokes and receiving the character-based output display. In the 1980s, the PC revolution started, which shifted much of the processing load to the desktop computer.

1. First Generation

The first generation of thin-client computers was a terminal-like computer. The servers performed the bulk of the processing duties, and consisted of the following three varieties:

- a. Basic Terminal—A variant thin client that was limited to a text-based display and keyboard input similar to the ‘dumb terminals’ of the mainframe computers. The basic terminal used the Microsoft Terminal Services protocols to access the server hosted applications (Esposito and Slack, 2009).
- b. Browser Terminal—A variant thin client that added a local Windows-like GUI, which resembled the desktop application GUI, but suffered from delayed reaction and webpage refresh. The Browser Terminal also supported the text-based command line

- along with the browser software to access the server hosted applications (Esposito and Slack, 2009).
- c. Line-of-Business (LOB) Terminal-A variant thin client that provided support for only a limited locally executed application. While the local application included a full GUI capability, the LOB terminal could not run any application, and still relied on much of the processing load to be handled by the server (Esposito and Slack, 2009).

In comparison with the conventional desktop PC, the first-generation thin clients offered the IT personnel who managed the organization's enterprise network greater user desktop control and security over their configuration and access. However, the loss of control by the end-user was not appreciated and contributed to an increase in user complaints.

The first generation of thin-client computers generally used low-power, low-capability processors, which cost less than the conventional desktop PCs. The thin-client had a small amount of Random Access Memory (RAM) and used a firmware-embedded operating system (OS), such as Microsoft Windows CE, Windows XPe, or an embedded variant of the Linux OS for communication. To communicate with the servers, the thin client used a set of protocols such as the Microsoft Terminal Services (MTS), Citrix Presentation Server (CPS), or X-Windows Display Manager Control Protocol (XDMCP) (Mui and Pearce, 1992). The thin-client user could then run shared multiuser applications connected to the application servers from within the thin-client GUI. Using the first generation of thin-client capabilities, either the individual applications could be published or an entire desktop could be published by the server to the thin-client device. Client Access Licenses (CAL) and any applicable user licenses for server applications increased the software cost in addition to the core server OS licensing fees. It can be argued that the total cost of ownership (TCO) for those first-generation thin clients was actually reduced compared to the conventional PC setup with the additional licensing fees. On the other hand, the TCO must also take into account the costs associated with size, weight, and power (SWaP). Finally, a critical problem with the first-generation thin-client technology was that all the users operated from one instance of the server OS. This shared use of the OS meant that if one user managed to lock up or crash the OS, all of the other users would stop operating until the effected OS was rebooted (Esposito and

Slack, 2009). From both a performance perspective and the cost of productivity downtime, the first-generation thin-client technology had to improve in order to become a viable solution.

2. Second Generation

Improving on the first generation of thin-client technology, the second generation aimed to satisfy the user needs while still satisfying the IT department's hardware centralization goal. At first, the solution was to create massive server farms with low-cost, rack-mounted commodity servers and load balancing software, which reduced the impact of OS crashes. The server farm model was then improved by employing server blade computers, which are fully-fledged computers that are integrated on a single board with circuit-based Input/Output (I/O) connectors plugged into a bus in a chassis. The difference between a blade computer and a single board computer (SBC) is that the SBC does not include peripheral I/O devices on the board such as a hard-drive, and therefore requires external cables. A blade computer is a stripped down server computer with a modular design optimized to minimize the use of physical space and energy. At the same time, the thin client became a commodity hardware item with lower cost and higher performance, which stressed the demands for network bandwidth using the same protocols as the first generation of thin clients. Unlike in the case of the first generation of thin clients, the user had a dedicated instance of his familiar Windows OS or Linux OS, which preserved the user's expected PC desktop experience. The second generation eliminated the situation where one user could take down the others in the case of an OS lockup or crash. While it can be argued that the TCO for a second-generation setup of thin-clients and blade servers costs somewhat less than does the traditional standard PC, it is less than ideal relative to the high cost of individual blade computers and their enclosure chassis. Another major cost contributor to the TCO of the second-generation thin-client solution is the continued requirement for desktop licenses (Windows or Linux) and for blade computer communication software for each user (Esposito and Slack, 2009).

3. Third Generation

The current third generation of thin-client technology aims to reduce the TCO through lower size, weight, and power (SWaP) of the first generation, while satisfying the same user and IT department's satisfaction goals of the second generation. To achieve a lower SWaP, the number of second-generation blade computers is being reduced through the implementation of virtual machine software. A virtual machine is not a new concept, but it has been improved through a new implementation that uses a Hypervisor, which is a modified OS that acts as a host for the additional virtual machines. Unlike the past virtualization technique of emulating the hardware in software, which suffered a performance hit, the Hypervisor solution merely acts as a resource allocator and scheduler for the hosted virtual machines that run directly on the hardware without any emulation. Although the Hypervisor solution requires addition resources, especially RAM to maintain performance, the lower commodity cost of higher density RAM makes virtualization a viable solution. There are three major virtualization solutions that meet enterprise network requirements: VMware server, Citrix's XEN, and the Linux based Kernel Virtual Machine (KVM). At the thin-client side, the latest Pano Logic and Sun Microsystems' SunRay 2FS thin-client technology use stateless devices, which means that the user session is running entirely on the server independent of the thin client, which is often referred to as truly thin-client or a "zero" thin-client (Esposito and Slack, 2009). There is no local storage device within the zero thin clients, and no processing takes place locally which enhances the Information Assurance (IA) posture of the enterprise network (meaning not exposing the thin-client to viruses, worms, malware, key-stroke loggers, etc.). Another IA advantage of not storing any data in the zero thin clients is that the same zero thin-client devices can be used for accessing multi-classification data and multi-compartmented networks by assuring there are no remnants of one session that can transgress to another.

C. CHAPTER SUMMARY

In this chapter, the theory of centralized servers and thin-client computing is introduced and described for the purpose of establishing a common understanding of how

the thin-client technology impacts the operations of an enterprise network. The chapter starts with a definition for client-server computing over a network using sockets and the IP protocol. A historical review of the thin-client and server advancements is provided as a means to gain understanding of the types of issues that came up and the lessons learned. The issues provide the starting point for the identification of the key factors that impact the operational capability of an enterprise network, which leads into the discussion on key architecture evaluation factors in Chapter III.

III. KEY ARCHITECTURE EVALUATION FACTORS

This chapter provides an overview of the key architecture factors that are required to successfully operate ONE-NET. The enterprise network operations complexities are described, which help in identifying the technology and labor cost factors due to the skill level and labor hours needed for a successful implementation. A detailed discussion is provided on each of the contributing factors for: Infrastructure and Network Services, Software Delivery and Management, Information Assurance, and Service Level Performance. Closing this chapter is a summary that includes a statement about the feasibility of the proposed thin-client solution.

A. OVERVIEW

Chapter II describes the technology aspects of the selected thin-client and server implementations over time and the general impact each generation had on the operations of an enterprise network. This chapter answers the first detailed question: Will replacement of the current thick-client desktops with thin-client desktops meet the operational requirements of ONE-NET?

In this chapter, the key factors that impact the enterprise network operation are evaluated to provide the basis for the following chapter's specific TCO analysis of alternatives. This chapter identifies four main categories of general enterprise network technology implementation disciplines: Infrastructure and Network Services, Software Deliver and Management, Information Assurance, and Service Level Performance. These categories cover the majority of the associated technology requirements for implementing and sustaining an enterprise network, and therefore have the greatest impact on the TCO.

B. INFRASTRUCTURE AND NETWORK SERVICES

At the highest level, an enterprise network consists of major facilities such as data centers and network infrastructure resources that connect the data centers with end-users.

In many cases, the network connectivity is achieved through lines that are leased from a third party or via satellite links when physical lines are not available (as in the case of some U.S. Navy OCONUS).

1. Footprint and Facilities

Deploying a thin-client based architecture relies heavily on shifting the computer processing power to the data center. Power, space and cooling constraints are the key factors that have a significant impact on the viability of the data center implementation. The BLII ONE-NET infrastructure installation commenced in fiscal year 2001 and concluded in early fiscal year 2004. Over this period, progress was made in incorporating the installation of the outside plant cabling (OSP), inside plant cabling (ISP), and data centers, including server farms, enterprise management systems, and information assurance (IA) suites (SSC PAC, 2010a).

During the second quarter of fiscal year 2004, the server farms at each data center throughout the ONE-NET enterprise network were enhanced to provide more processing power and larger storage space to meet the growing fleet operational requirements. In January 2005, the desktop refresh and user migration effort had moved all the customers to the upgraded ONE-NET infrastructure. As of today, ONE-NET provides IT and critical telecommunications support services to approximately 33,000 workforce of Navy uniformed and civilian members across the 14 sites within Europe, Far East, and the island nation of Bahrain. The desired end state of ONE-NET is to deliver the enterprise network capabilities required for effective Command and Control (C2) of all the warfighter and business missions carried out by the U.S. Navy overseas commands. The goal is to sustain and normalize day-to-day operations and maintenance across all ONE-NET sites, while replacing all the legacy networks in alignment with the DON CIO vision for the future NGEN capabilities and integration (SSC PAC, 2010a).

The data center solution for the TNOSC and LNSC is built upon several server building blocks for data processing and storage as shown in Figure 2.




Figure 2. TNOSC and LNSC Unclassified Server Farm Layout (From CSC, 2009)

Under the current lifecycle upgrade, the number of servers and storage devices at each data center located in the TNOSCs and LNSCs will be reduced by more than 50%, which is a key factor in lowering hardware costs as well as heating, ventilation, & air conditioning (HVAC) costs. A total of 320 physical devices will be virtualized at the data centers throughout the fourteen TNOSC and LNSC sites. The new servers, storage, and backup devices will be installed using existing power, HVAC infrastructure. The server farm refresh solution replaces Dell PowerEdge 2650, 2850 and 2950 models with Dell's latest 11th generation hardware. The required servers are limited to two models, the R610 and the R710. The R610 is a 1 rack-unit (RU) server that supports up to two 64-bit capable quad-core Xeon 5500 processors and 96GB of RAM. A RU is a standard measuring unit for IT racks where 1 RU is equal to 1.75" (4.45 cm) in height, and the significance of this measurement unit has to do with how densely a standard rack can be populated. These servers will be used for non-virtualized services such as domain controllers. The R710 is a 2 RU server that has been designed with hardware support for virtualization. The R710 has 125% more memory and more integrated I/O than the previous generation servers, and includes the embedded Hypervisor technology, all in an energy efficient, low profile design. These servers support up to two 64-bit capable quad-core Xeon 5500 processors and 144GB of RAM making them optimal for

virtualization environments. The importance of this data is primarily to show that the current lifecycle upgrade investment can be reused to support the proposed deployment of thin-clients. The Dell EqualLogic PS Series will be used for the storage area network (SAN). The SAN creates a virtualized Internet Small Computer System Interface (iSCSI) SAN that is capable of supporting a mix of drive speeds and capacity within the same peer group and work together to automatically manage data, load balance across resources, and expand to meet growing storage needs. The PS Series arrays can be easily added to the existing EqualLogic storage infrastructure as modular building blocks for a future SAN expansion (CSC, 2009).

The logical architecture for a TNOSC data center is shown in Figure 3 as a visualization aid in understanding how the specific hardware used for the TCO calculations is integrated.

The top of Figure 3 shows the new 10 gigabit per second Ethernet (10GB) that provides the large bandwidth connection to the Virtualization Servers that the thin-clients access via the Virtual Switch Network. The applications in the five lower blue ovals and the one gigabit per second Ethernet (1GB) are the existing core services that are connected to the new Virtualization Servers via a redundant set of Cisco Nexus switches.




Figure 3. High-Level TNOSC Data Center Infrastructures (From CSC, 2009)

To extend the server farm refresh's core virtualization infrastructure to support the proposed thin-client solution, each TNOSC and LNSC will need to be able to grow through adding building blocks with additional servers and SAN disk arrays. Figure 4 shows an example building block that supports 1,000 virtual machine desktops. As shown in Figure 4, the building block consists of a layered implementation on top of the physical hardware using a VMware VCenter Server that supports two clustered hosts and shared storage.




Figure 4. VMware® Building Block for 1000 Virtual Desktop Users (From VMware, 2010)

Desktop virtual machines for end users do not require as much disk space and processing resources. The amount of system disk space required depends on the number of applications required in the base desktop image, where a typical applications suite includes Microsoft® Word®, Excel®, PowerPoint®, Adobe Reader®, Internet Explorer®, McAfee® Antivirus, and PKZIP. The amount of disk space required for user data depends on the role of the end user and organizational policies for data storage. Although support for the thin-client capability requires additional virtual servers and SAN capacity, these additional resources will have minimal impact on the existing power, cooling, and available space.

2. Networking Resources

Thin-client architecture relies on a reliable distributed network-computing infrastructure, where operation can be severely impacted by an inability to connect to the central server resources. Both latency (the time it takes a data packet to travel from the source sending it to the destination receiving it via the network) and bandwidth (the amount of data that can travel over the network within a given period of time) are important factors. In the proposed thin-client computing scenario, the remote servers generate everything a user sees on their screen. High latency has a serious impact on the perceived response of the system, while low bandwidth affects the time it takes to get large chunks of data like bitmaps to the user's screen. Additionally, if transport encryption protocols are being used, the latency will be increased even further.

Therefore, the user experience depends very much on the network infrastructure resources' ability to provide sufficient bandwidth. It is important to note that while the satellite data links have limited bandwidth, the critical links for supporting the thin-client solution between the local data centers and their users are not over satellite links. Therefore, the main resourcing cost involves the switching and routing devices in the data center.

The current server farm refresh uses Cisco Nexus 5010 and Nexus 2148T switches to provide ONE-NET with auto-negotiating 10/100/1000/10000 mega bits per second (Mbps) throughput capabilities and an unified low-latency network fabric (CSC, 2009). In addition, the Cisco Nexus 1000V virtual switches are installed inside the virtual server environments to provide a direct extension of the Nexus networking capabilities to the virtual machine level. These software switches operate inside the VMware virtual environment and provide a direct extension of the network fabric. To separate network traffic for performance and security, six separate Virtual LAN (VLAN) networks are created as shown in Figure 5. Figure 5 shows the notional view of the same data center architecture described in Figure 3, but from the VLANs' overlay perspective. The overlaid VLANs are described in Figure 5 by the blue clouds below the Internet cloud. Each of the VLANs is served by a combination of network data traffic from the Windows core servers, virtual servers, backup, and data storage. Expanding VLAN capacity in the future is done through the configuration of the Cisco's virtual switch technology.




Figure 5. Networking Server Farm with Six VLANs (From CSC, 2009)

To optimize network performance, the network infrastructure will take advantage of application-validated protocol-specific acceleration technology. The Cisco Wide Area Application Services (WAAS) provides application-specific acceleration features for both encrypted and non-encrypted applications. This technology can improve network performance over the WAN and reduce the effects of latency and bandwidth utilization, which provides a user experience similar to using a desktop PC. This is done through protocol acceleration, read-ahead, operation batching, multiplexing, and safe caching. Applications include Microsoft file services (Common Internet File System [CIFS]) and Microsoft Exchange (Messaging Application Programming Interface Remote Procedure Call [MAPI-RPC]), plus numerous other application protocols. The Cisco WAAS Replication Accelerator mode is optimized for WAN links from 20 Mbps up to OC-12

(622 Mbps) and a latency greater than 20 milliseconds. In addition to new core switches, the WAAS accelerators can optimize the current ONE-NET WAN bandwidth. However, in the event that ONE-NET significantly increases the bandwidth in the future, it is likely that the WAAS devices will need to be upgraded or replaced to support the additional load.

C. SOFTWARE DELIVERY AND MANAGEMENT

Installing and managing software on a PC is relatively simple when handling only a small number of PCs, but the problem becomes a lot more challenging when dealing with tens of thousands of computers or more. It is challenging because, on one hand, the enterprise level organization cannot afford to let each user manage his or her own PC because of licensing constraints and security concerns, and on the other hand, centralized delivery and management adds additional complexity and costs. The current ONE-NET refresh plan is intended to continue to support PC devices for each user and, as of March 2009, over 6400 PC refresh seats have been completed on nine of the fourteen sites (SSC PAC, 2010b). Although deploying the proposed thin-client solution requires that the data centers be upgraded first, the current plan calls for the completion of the data centers refresh by the second quarter of fiscal year 2012 (SSC PAC, 2010c). According to the ONE-NET plan for completing the PC deployment and the PC refresh cycle, there will be sufficient time to deploy the thin-client solution after the data center refresh completion date.

1. Baseline Configuration Management

Currently, desktop delivery on ONE-NET is done primarily with thick-client PCs at the user end and servers at the data center using the client-server architecture. ONE-NET provides a common desktop software configuration that includes all the Approved Product List (APL) applications, in what is known as the Workstation Baseline Software Configuration (WBSC), to each PC in the enterprise. The APL of desktop applications, operating system patches, application software updates, and hot-fixes for Information Assurance Vulnerability Alerts (IAVA) make up the ONE-NET baseline desktop PC software configuration. The baseline PC hardware configuration is Dell computer-

specific, where the WBSC baseline build can be applied to any Dell model that is supported by Dell X-Image technology. The intent is to have a single standardized baseline for all the Dell PCs, notebooks and tablets throughout ONE-NET. The WBSC image build solution is intended for use by all ONE-NET users and administrators when reimaging existing PCs, or when applying to new PCs, and it applies to both the classified and unclassified network domains. The current WBSC image build version is 118, which is based on Windows XP Service Pack 2 and incorporates all baseline updates, security settings, enterprise configurations, legacy compatibility, and functional updates. This build is compatible with all desktop platforms going back at least five years. Table 1 lists the software components required to design and build the initial WBSC image. The software components are the Dell X-Image, PowerQuest and the VMware. The Dell X-Image is used to incorporate multiple Dell PC platforms' drivers into the image. The PowerQuest is used to create a bootable optical disk format after completion of the image. The VMware is used to create a virtual image from the initial image build (Lelfwitch, 2009).

Table 1. Required Software Components to Create WBSC Image (From Lelfwitch, 2009)

Software Application	Developer	Version	Description/Function
X-Image	Dell	3.01	Incorporates drivers for multiple platforms into the image
PowerQuest	Symantec	5.51	Utilized in creation of bootable DVD set
VMware	VMware	6.5	Build environment

2. Image Build Process and Administration

An overall view of the process for building a new WBSC image baseline is shown in Figure 6. This is a continuous process of modifying previous images with a strategy around baseline configuration management. Due to the many different security modifications and vulnerability patches, ONE-NET uses the Windows Patch Management System (WPMS) to periodically apply the PCs through Tivoli Enterprise Management System (EMS). The current WBSC version 118 Image Build also supports virtual machine (VM) workstation and the following platforms.

- Desktops: OptiplexGX270, OptiplexGX280, OptiplexGX620, Optiplex740, Optiplex745, Optiplex755, Optiplex760.
- Notebooks/Tablet: Latitude D400, Latitude D410, Latitude D420, Latitude D430, Latitude D600, Latitude D610, Latitude D620, Latitude D630, Latitude E6400, InMotionLE1600 (Tablet), Latitude XT (Tablet) (Lelfwitch, 2009).




Figure 6. WBSC118 Image Build Architecture (From Lelfwitch, 2009)

The image build process includes inputs from update packs (UP), which is done in parallel with ONE-NET EMS, and pushes the software updates to PCs in order to maintain a stable baseline that is consistent with previous versions of the WBSC image. Both update packs and Tivoli patch management pushes to PCs contain the necessary security and functional updates from the monthly Microsoft patch releases in response to the Joint Task Force-Global Network Operations (JTF-GNO) Information Assurance Vulnerability Management (IAVM) notifications. IAVM compliance is required in order to maintain a secure WBSC ONE-NET baseline build. The build also contains driver cab

files for all desktop and laptop systems, including both current and legacy support for at least five years (Lelfwitch, 2009). While Figure 6 described the specific WBSC Image 118 build process, there are three possible update processes for building a WBSC image 118, as shown by the three progression bars of Figure 7 (Figure 6 is represented by the middle progression bar). The top progression bar shows a WBSC image 118 build from a WBSC image version 117 with XP Service Pack 2 build and additions of update packs. The update packs include UP117A, UP117B, UP117C and UP118. Within these update packs are the Microsoft security patches, third-party application updates and IAVA updates. The bottom progression bar shows a WBSC image 118 that started with an older version of WBSC images 115 or 116 or 117, and all required accumulative updates that can be delivered by using the Tivoli Enterprise Management Service (EMS).




Figure 7. Optional WBSC Baseline Image Build Progressions (From Lelfwitch, 2009)

3. Application Management and Licensing

In order to calculate the TCO impact of application management and licensing, the WBSC baseline image is identified, which includes the Windows Operating System (OS) and main components, Microsoft Office Professional 2003 Service Pack 2, file viewer applications, multi-media player applications, utility applications, security applications, enterprise management client and miscellaneous applications as shown in the multi sectional Table 2. Some applications have free software licenses and some

applications require paid licenses. The license information in Table 2 shows how the licenses are applied to the WBSC. The licenses for Windows OS components and Windows Microsoft Office Professional Suit are covered by the ONE-NET Microsoft enterprise license agreement contract. The 3-year PC refresh takes into account that these Microsoft licenses are paid for once per refresh.

Table 2. Baseline Applications on WBSC Image Build (SSC PAC, 2009a)

Baseline Operating System and OS Main Components

Manufacturer	Application Name	Version	Build / Version	Licence Info
Microsoft	Windows XP Professional	5.1.2600 (SP2)	2600.xpxp.050622-3462: SP2	Part of OS
	Internet Explorer	6	6.0.2900.2180.xpsp_sp2_gdr.080814-1233	Part of OS
	Outlook Express	6	6.00.2900.2180.(xpsp_sp2_rtm. 040803-2158	Part of OS
	Windows Messenger	4.7	4.7.0.3001	Part of OS
	DirectX	9.0C	6.5.2600.3367	Part of OS
	Windows Media Player	10	10.00.00.4058	Part of OS
	MDAC	2.8 SP-1	2.81.1117.6	Part of OS
	Java Virtual Machine (MSJVM)	5	5.00.3810.0	Disabled, replaced by Sun Java
	.Net Framework	1.1	1.1.4322.573	Part of OS
	.Net Framework	2.0	2.0.50727.842	Part of OS
	Remote Desktop Client Update	5.2	5.2.3790.0	Part of OS
	OWA S/Mime Control	6.5	v6.5.7226.0	Part of OS
	MS XML Parser	1.0 SP1	4.72.3110.0	Part of OS
	MS XML Parser	2.6 SP3	8.30.9529.0	Part of OS
	MS XML Parser	3.0 SP9	8.90.1101.0	Part of OS
	MS XML Parser	4.0 SP2	4.20.9870.0	Part of OS
	MS XML Parser	5.0	5.20.1087.0	Part of OS
	MS XML Parser	6.0 SP2	6.20.1099.0	Part of OS
	MS Language IME for Japanese, Chinese and Korean	2002a	IME 2002	Part of OS

Microsoft Office Professional 2003 (Service Pack 2)

Manufacturer	Application	Version	Build / Version	Licence Info
Microsoft	Access 2003	2003	(11.8166.8221) SP3	ONE-NET Enterprise Licence
	Excel 2003	2003	(11.8231.8221) SP3	ONE-NET Enterprise Licence
	InfoPath 2003	2003	(11.8165.8221) SP3	ONE-NET Enterprise License
	Outlook 2003	2003	(11.8217.8221) SP3	ONE-NET Enterprise License
	PowerPoint 2003	2003	(11.8212.8221) SP3	ONE-NET Enterprise License
	Publisher 2003	2003	(11.8212.8221) SP3	ONE-NET Enterprise License
	Word 2003	2003	(11.8227.8221) SP3	ONE-NET Enterprise License
	MS Office 2007 File Converters	2003	12.0.6320.5000	ONE-NET Enterprise License

Misc Applications

Manufacture	Application	Version	Build / Version	License Info
DB Web Sign	DB WebSign Browser Plug-In (For DTS)	2.3	2.3.1.12	Free
Sun Java	Sun Java Runtime Env (JRE)	1.6x	v1.6.0_07-b06	Free
SOCOM	Install SOCOM's Outlook Email Classification Add-In	2.x	v2.2	free GOTS app designed by SOCOM

Enterprise Management Client

Manufacture	Application	Version	Build / Version	License Info
IBM	Tivoli Management Agent (TMA) - Endpoint Manager Client for software distribution.	4.1x	4.1.138	ONE-NET Enterprise License

File Viewer Applications

Manufacture	Application	Version	Build / Version	License Info
Adobe	Reader (Full) with Font packs for Japanese, Chinese, Korean and European	9	9.0.0	Free
Autodesk	Volo View Express	2.01	811	Free
Microsoft	Visio Viewer 2003	11	11.0.3709.5614	Free
	Windows Journal Viewer	5.1	2600.xpsp_sp2_gdr.080814-1233	Free

Utility Applications

Manufacture	Application	Version	Build / Version	License Info
WinZip Computing	Winzip	9.0 SR-1	6224	One time purchase of 15,000 seats
Symantec	AntiVirus Corporate Edition	10.1x	10.1.6.6010	DOD Enterprise License
Sonic	Roxio Easy CD Creator	5.3.5.17	A00	Licensed for each Dell PC

Multi-Media Players

Manufacture	Application	Version	Build / Version	License Info
Apple	QuickTime Player	7.55	7.5.0	Free
Adobe	Macromedia Shockwave Browser Plug-in	10	10.1.0.11	Free
	Macromedia Flash Player	9	9.0.124.0	Free
	Macromedia Authorware Web Player		4.0.0.70	Free
InterVideo	InterVideo WinDVD	4.0	4.0, DXVA B11.052C13.10325.0000D 00000	Licensed for each Dell PC

Security Applications

Manufacture	Application	Version	Build / Version	License Info
DOD	Root CA PKI Certificates	3.06A		Free
Tumbleweed	Desktop Validator (For CLO)	4.72	4.72	SPAWAR Seat License
MS CAPICOM	CAPICOM Module	2.1x	2.1.0.2	Free
ActivCard	Active Client for CAC - PKI	6.1		SPAWAR Seat License
Juniper Networks	NetScreen Remote VPN Client	8.5	10.3.5 build 6	SPAWAR Seat License

D. ENTERPRISE DEPLOYMENT AND INFORMATION ASSURANCE

ONE-NET currently uses the Altiris enterprise deployment solution, which provides an automated process for collecting data and maintaining a configuration management across each enclave for every seat deployment to the latest approved baseline configuration. An enclave is a grouping of networked devices that share similar classification authority levels. A “seat” is the collection of resources required to support an individual user. Along with the Altiris enterprise deployment solution, the PCs are first placed on the staging VLAN, which offers a segregated area where all the PCs can reside until they conform to the configuration and security requirements. Upon validation, the PCs can then be deployed to the ONE-NET enclave. The architecture of the Altiris enterprise deployment system contains the connections and interfaces as shown in Figure 8 (Martinez, 2010). Figure 8 shows the enterprise deployment hierarchy where the top level global notification server (in red) is located at SPAWAR Systems Center Pacific, San Diego. This server replicates security policies, image packages and other configuration items to the regional notification servers located at the TNOSCs in: Yokosuka, Japan; Naples, Italy; and the island nation of Bahrain. The regional notification server at each TNOSC (in blue) then replicates the same data to LNSCs (in green) within the region. At each LNSC, the local notification server applies security policies, image packages and configuration items to the corresponding end user PC system. The inventory data is reported to the local notification servers using the Symantec Management Agent.




Figure 8. Altiris® Enterprise Deployment Architecture (From Martinez, 2010)

The Altiris enterprise deployment solution consists of the Symantec Management Platform version 7.0 SP4, Deployment Solution version 7.1, SQL, IIS, and Symantec Management Agent (Martinez, 2010).

The ONE-NET's enterprise desktop imaging deployment provides a centralized deployment management that consists of the following capabilities:

- Standard jobs, scripts, images and packages from a central location.
 - Automated replication of jobs, scripts, images and packages.
 - Centralize imaging deployment and maintaining desktop configuration control.

A staging VLAN provides a quarantined area for PC installation denoted by the Workstations icon in the bottom left of Figure 9. The quarantine is done by implementing a Level-1 network transport architecture as shown in Figure 9, using a layer 2 IP network VLAN with defined access control lists (ACLs) to isolate the new PCs

from the targeted ONE-NET enclave on the right. The ACLs provide the PCs with limited access to ONE-NET resources that are used to update noncompliant PCs (McDaniel & Falcone, 2008). Figure 9 also shows the connectivity to ONE-NET's Tivoli EMS and Symantec resources via the core distribution switch, which allows the PCs to receive the latest patches and security updates. Connectivity to the existing Information Assurance (IA) suite Retina server allows the PCs to be validated for security compliance before being sent and installed in the ONE-NET enclave (Hanada, 2009).




Figure 9. Staging VLAN Architecture (From Hanada, 2009)

The process leading to the deployment of the PCs is described in Figure 10, where the PCs go through two major steps for imaging and staging. The process starts with imaging the PCs and takes about 1.5 hours to complete. After the PCs are imaged, the first step is completed with an initial Retina security scan (RSS) before placing the PCs in the quarantined staging VLAN. A second step of seat to application mapping (STAM) process is then applied to the PCs, where the latest software patches are installed (pushed) on the PCs from the Tivoli EMS. Completing the process is a final RSS prior to the deployment of the validated PCs on the network.




Figure 10. Post Imaging PC Patching Process (From Graham, 2009)

The performance required for this process is as follows (Graham, 2009).

1. Baseline Imaging: 1.5 hours
2. Retina Scan on PC: 10 minutes
3. Post baseline image update Seat to App Mapping (STAM): 1 hour
4. Application patches using Tivoli Windows Patch Management System (WPMS): 1 hour (maximum)
 - a. The ability to manually initiate PC Inventory using WPMS
 - b. The ability for manual or automatic start of the remediation (patch) process using WPMS
 - c. Completion status notification of workstation remediation by WPMS
5. Retina scan post remediation: 10–20 minutes

The requirement timeline as defined above supports a PC deployment of 32 desktop systems in approximately a 4-hour period (Graham, 2009).

Similar to the deploying process, the ONE-NET Tivoli WPMS also provides a capability for facilitating a rapid delivery of updates to non-compliant PCs that includes a robust reporting system for compliance statistic. The Tivoli WPMS lifecycle begins on the second Tuesday of every month when Microsoft posts its updated security bulletin, and lasts through the rest of the month until Microsoft releases the next security bulletin (Shimoko, 2008a). The automated inventory system is an integral part of and a mandatory prerequisite for the WPMS. Before WPMS non-compliant PCs and the patches they need are identified, an inventory scan is done on the PCs. There are 3 different scan configurations that are utilized on ONE-NET: Initial Nightly Scan, Weekly Difference Scan, and Periodic Full Scan (Shimoko, 2008b).

The WPMS process shown in Figure 11 applies to both desktop PCs and Windows servers. The workflow process for testing, packaging, deploying and reporting monthly patches and updates covers 26 detailed steps. The top blue ovals in Figure 11 describe the overall two parallel activities, one for the Microsoft patch releases and the second for the information assurance vulnerability management (IAVM) notifications. The workflow diagram shows a process for mapping the JTF-GNO information assurance vulnerability alerts (IAVA) release number with any applicable Microsoft security vulnerabilities. The applicable patches are then manually tested prior to installation in order to ensure the WBSC functionality is maintained. After testing is completed, the patches are released to get approval through the approval procedure prior packaging and deployment. The final step completing the process in Figure 11 ends with reporting the updated status to the DoD's Vulnerability Online Compliant Report System (OCRS) for patch compliant.




Figure 11. ONE-NET Windows Patch Management Process (From SSC PAC, 2009b)

E. SERVICE LEVEL PERFORMANCE

Current ONE-NET network infrastructure and services have demonstrated the ability to meet the target service availability of 99.5% based on NNWC's Service Level Objectives requirements (NNWC, 2001). Based on the review and assessment discussed in Section B, the current network latency and bandwidth are expected to support the thin-client implementation using a virtual desktop delivery solution without impacting the current Quality of Service (QoS). The network infrastructure within the data center provides a robust high-speed Local Area Network (LAN) with 10 gigabits per second (10GB) Ethernet backbone service as the maximum utilization available to the proposed thin-client solution. The thin clients can use the existing networking infrastructure to implement the PC over IP (PCoIP) display protocol from VMware, which provides an optimal display performance across LANs and WANs. The PCoIP is the default protocol in the Windows or Linux software that connects servers through View Manager to their thin-client devices. The PCoIP also supports productivity applications like Microsoft Office and rich media like video, flash, and graphics (VMware, 2010).

As shown in Section B, subsection 2 (Networking Resources), the combination of the Cisco unified low-latency switches, virtual switches, and WAAS, offer an optimized network performance that is capable of supporting the proposed thin-client solution with expected similar existing QoS. The technology greatly reduces the complexity and cost of the current PC deployment and updating process because it will all be localized in the data center. The PC deployment and updating process localized in the data center is a significant TCO benefit in terms of staffing reduction that will impact the TCO calculations in Chapter IV.

F. CHAPTER SUMMARY

By reviewing the current ONE-NET architecture, this chapter covers the various aspects that impact costs of running and managing data center facilities, network resources, software delivery, updating, patch management, and security. A final look at the service level performance of the current infrastructure resources provides the assurance that the thin-client proposal is at least technically feasible and positively

answers the first detailed question: Will replacement of the current thick-client desktops with thin-client desktops meet the operational requirements of ONE-NET? In the next chapter, the actual TCO costs are calculated and compared between the current thick-client and thin-client alternative solutions.

IV. TOTAL COST OF OWNERSHIP

This chapter introduces the concept of TCO and identifies the inputs needed for the cost per seat calculation model as well as for the power and cooling cost tool. Using the identified inputs, the TCO is calculated for the current thick-client implementation and then for the proposed thin-client implementation. The resulting TCOs are then used in the analysis of and comparison between the two alternative implementations of ONE-NET.

A. TOTAL COST OF OWNERSHIP CONCEPT

The Total Cost of Ownership (TCO) is a financial estimate concept to support acquisition and planning decisions for a wide range of assets that incur significant maintenance and/or operating costs over a long period of use. It quantifies and measures the various associated costs for the purpose of planning future budgets, and offers excellent insight into business processes and levels of service which will help improve efficiencies and effectiveness. The object of good business is to attain the lowest TCO (Solution Matrix, 2010).

In ONE-NET, the TCO calculations are used for both gauging the required capital investment by the ONE-NET program office and gaining insight to various cost aspects through the analysis process. The insights gained can then be used to identify potential means for reducing the TCO, which is a primary management objective. Chapter III describes four key architecture evaluation factors that have an impact on the enterprise network operations, and proves that the thin-client desktop proposal is technically sound. This chapter delves into the TCO calculations in order to answer the second detailed question: Will the TCO of ONE-NET over the next seven years be low enough to justify the additional risk and up-front investment in deploying the thin-client desktops?

The four major TCO components are labor, hardware, software and transport. The following section describes the major cost components and assumptions used in the TCO calculations.

B. MAJOR COST COMPONENTS AND ASSUMPTIONS

Calculating the TCO for ONE-NET requires a major data gathering, which cannot be done without some initial assumptions that bound the effort. The ONE-NET cost per seat (CPS) model was developed in year 2005 by NNWC and is used for the TCO calculations. As described earlier, the TCO consists of the four major costs components—labor, hardware, software and transport. The assumptions for each of the major cost components are now described.

1. Labor Costs

Labor costs are defined as the burdened costs of all dedicated staff required to operate and maintain ONE-NET (NNWC, 2005). The staffing model used for this analysis is based on the full Future Mode of Operations (FMO), which considers labor for different types of personnel such as civilian service, military, foreign nationals and contractors. The FMO labor costs were compiled for the expected labor pool required to support a 27,284-seat enterprise network and maintain up to 41,000 users (NNWC, 2005).

2. Hardware Costs

The CPS model costs include an estimate of the total number of PCs (one PC equals one seat). There are additional hardware costs for the data centers, but they do not include the costs of all servers, switches, and routers because most are reused. Other hardware costs relevant to the analysis include resources used by non-seat users such as help desk and engineering support. The analysis is based on a three-year refresh cycle duration with a PC count growth tied to the expected seat count. However, since the seat count is based on the fixed FMO number of 27,284 seats with one-third of the seats refreshed each year (approximately 9,095 PCs), the growth is actually constant over the three year period (NNWC, 2005).

3. Software Costs

Client Access Licenses (CALs) are associated with each seat. When a PC is purchased, it includes the CALs in the software package that are good for three years (for

the CPS model purpose, the CALs are accounted for in the first year of the PC purchases and then zeroed for the next two years). After three years, the CALs are accounted for annually and can be funded separately via the Program Objective Memorandum (POM) process, which is independent of the initial acquisition funding. Although the follow-up POM funding is not guaranteed, the CAL costs for year four and beyond are still included for the TCO calculation (NNWC, 2005).

4. Transport Costs

Network transport costs are based on the Defense Information Systems Agency (DISA) Enhanced Planning Process (EPP). All DISA circuits are priced according to the Defense Working Capital Fund (DWCF) “Telecommunications Pricing Guide.” Channel Service Unit (CSU)/Data Service Unit (DSU) costs are in accordance with the DWCF. A charge is incurred each time the bandwidth reaches a new circuit size (NNWC, 2005). Non-DISA point-to-point circuit CSU/DSU costs vary widely (by country, by location, by commercial entity) and their exact amounts are not known, so the DWCF Telecommunications Pricing Guide is used as an average CSU/DSU cost figure (NNWC, 2005). The Navy does not have a standard for metrics in terms of bandwidth per user (e.g., number of users per location, cost per Megabit). Therefore, bandwidth growth is based on the following distribution: (1) No growth or no change in the first year; (2) DISA circuit growth rates of 10% for the following four years and 5% for two more years; and (3) circuit bandwidth remains static in the future years.

C. COST ANALYSIS AND TCO COMPARISON

Based on the assumptions described in the previous sections, each major cost component is formulated in the cost per seat (CPS) model as follows.

1. Labor Cost = LNSC Labor Cost + TNOSC Labor Allocation + Total Other Labor Costs, where LNSC Labor Cost is defined as the cost associated with the required number of staff personnel at the seat locations to provide operational and maintenance support; TNOSC Labor Allocation is defined as the cost of TNOSC staff personnel allocated to support the LNSC service desk; and Total Other Labor Costs consist of training, travel and consumables costs for staff personnel.

2. Hardware Cost = Total Cost for Unclassified PCs + Total Cost for Classified PCs + Total Cost for Notebooks.
3. Software Cost = Total Cost for Number of CALs required for the 27,284-seat network.
4. Transport Cost = Total Cost for the Number of DISA Circuits and Non-DISA Circuits throughout the enterprise.

The ONE-NET cost per seat (CPS) model was initially developed by NNWC for the period of fiscal year 2005 (FY 05) through FY 11, so it provided a good starting point for the TCO calculations covered by this thesis. The cost model used the projected inflation index rates given in Table 3 (NNWC, 2005).

Table 3. Original Inflation Index (From NNWC, 2005)

	FY 05	FY 06	FY 07	FY 08	FY 09
<i>Inflation Index</i>	1.000	1.012	1.021	1.030	1.040

Based on Table 3 from the CPS model, one dollar in FY 05 will be worth \$1.107 in FY 09. The estimated inflation rate over the 5-year period is thus 10.7%.

Since the thesis is concerned with FY 10 to FY 16, the CPS model costs were adjusted for the past inflation by 10.7% starting in FY 10. A similar inflation index distribution was used by the CPS model to extend the calculation to FY 16 and applied to the raw data results of the model, which are provided in Appendix A.

1. Thick-Client Cost Analysis

The thick-client calculations for the four major cost components are based on the current adjusted inflation index rate and are shown in Table 4. Note that the first three years do not have software costs as software is already included with the purchased PCs.

Table 4. Four Major Cost Components for Thick-Client Solution

Major Cost Components	FY 10 (Million)	FY 11 (Million)	FY 12 (Million)	FY 13 (Million)	FY 14 (Million)	FY 15 (Million)	FY 16 (Million)	Total per Component
Labor	\$70.24	\$72.65	\$75.13	\$77.71	\$80.38	\$83.14	\$86.00	\$545.25
Hardware	\$13.00	\$13.50	\$13.57	\$13.84	\$14.06	\$14.29	\$14.70	\$96.96
Software	\$0.00	\$0.00	\$0.00	\$3.35	\$5.87	\$6.00	\$6.12	\$21.34
Transport	\$7.60	\$18.39	\$18.85	\$19.31	\$19.73	\$20.14	\$20.59	\$124.61
TOTAL	\$90.84	\$104.54	\$107.55	\$114.21	\$120.04	\$123.57	\$127.41	\$788.16

Based on the values from Table 4, the sum of the total thick-client costs for labor, hardware, software, and transport over the seven year period is about \$788 million.

To help in understanding how the values in Table 4 were calculated, the calculations for FY 13 are given below as an example. Since this thesis is concerned with the period of FY 10 through FY 16, the CPS model had to be adjusted from starting in FY 05 to starting in FY 10. This was done by applying the 10.7% adjustment to the output data from the CPS model. Therefore, to provide FY 13 as an example, the corresponding year from the CPS's original model is FY 08. The FY 08 excerpt data from the CPS model calculations for the thick-client solution is provided in Appendix A. Note that all the CPS model output data was rolled up by region to provide a more consolidated view, where the regions are the island nation of Bahrain, Far East, and Europe.

While Appendix B provides the full regional rollups for FY05 through FY11 that were done on the output data of the CPS model, for this example, the specific table for FY 08 is used to show the corresponding data rollup of Appendix A. Table 5 shows the specific rollup for FY08 taken from Appendix B. The data in Table 5 shows the rolled up data for the four cost components prior to the 10.7% adjustment to FY 13.

Table 5. FY08 Regional Rollup for Thick-Client Solution (Appendix B)

FY 08	Labor	Hardware	Software	Transport
Bahrain	\$7,565,545	\$1,009,097	\$244,811	\$4,531,136
FE	\$38,023,250	\$8,194,600.88	\$1,989,904.28	\$7,001,710
EU	\$24,608,817	\$3,297,472.16	\$794,496.18	\$5,906,350
total	\$70,197,612	\$12,501,170.42	\$3,029,211	\$17,439,196.91

Applying the 10.7% adjustment to the entire data in Table 5 results in the data used in Table 4 for FY 13, where for example applying the adjustment to the total labor cost is as follows:

$$\text{Labor: } \$70,197,612 * 1.107 = \$77.71 \text{ million (rounded up)}$$

Applying the 10.7% adjustment to the total Hardware in Table 5 results in the following hardware cost for FY 13 in Table 4:

$$\text{Hardware: } \$12,501,170.42 * 1.107 = \$13.84 \text{ million (rounded up)}$$

Applying the 10.7% adjustment to the total Software in Table 5 results in the following Software value for FY 13 in Table 4:

$$\text{Software: } \$3,029,211 * 1.107 = \$3.35 \text{ million (rounded up)}$$

Applying the 10.7% adjustment to the total Transport in Table 5 results in the following Transport value for FY 13 in Table 4:

$$\text{Transport: } \$17,439,196.91 * 1.107 = \$19.31 \text{ million (rounded up)}$$

Therefore, the total for the four major cost components under FY 13 is:

$$\text{Total: } \$77.71 + \$13.84 + \$3.35 + \$19.31 = \$114.21 \text{ million}$$

Because one of the major benefits for the thin-client solution is the reduced power and cooling usage, the cost of power and cooling is added to the TCO calculations for both thick-client and thin-client solutions. The operating and cooling power is calculated by using a calculation tool provided by VMware® Inc. (VMware, 2010). The annual (360 days) desktop power and cooling costs for 27,284 thick-client seats are shown in Table 6. Therefore, the annual operating and cooling power costs for the current ONE-NET thick-client solution are provided in Table 6, where the total annual cost is computed as $21,687,506 \text{ kWh} * \$0.1016/\text{kWh} = \$2,203,451$.

Table 6. Annual Cost of Operating and Cooling Power on Thick-Client Solution
 (From VMware, 2010)

Desktop Power and Cooling Factors	Thick-Client Solution Values
Hardware operating power (W/h)	102
Air conditioning cooling power (W/h)	82
Operating hours per day	12
Power consumption per seat (W)	2,208
Total power consumption per day (W)	60,243,072 W
Total power and cooling (kW/yr)	21,687,506
Cost per kWh	\$0.1016
Total Annual Cost	\$2,203,451

Based on a fixed annual cost of \$0.1016 per kilowatt hour (kW h) over the span of the contract, the total seven year cost for the current ONE-NET thick-client solution is then \$15.4 million (rounded up).

2. Thin-Client Cost Analysis

Cost analysis for the thin-client uses similar calculations as for thick-client except for the hardware costs and labor input values into the CPS model. One of the advantages of the thin-client solution is that fewer staff personnel are required to manage the help-desk and administrative support because the software, configuration management, Information Assurance policy enforcement, etc., are all done in the data center. Furthermore, because the simplicity of the thin-client, there are fewer trouble calls from the users. Therefore, while the number of managers was kept the same, the number of service desk staff used in the labor portion of the cost per seat model is half the number of those used for the thick-client solution. Another advantage of the thin-client solution is that the hardware cost for client devices has a lower cost compared to the thick-client PC's cost. An unclassified PC, including monitors, costs \$1,216 and a classified PC, including monitors, costs \$1,324. A market estimate of \$450 is used in the CPS model as

an average cost for both classified and unclassified thin-client devices, including monitors. The thin-client calculations for the four major cost components are shown in Table 7.

Table 7. Four Major Cost Components for Thin-Client Solution

Major Cost Factors	FY 10 (Million)	FY 11 (Million)	FY 12 (Million)	FY 13 (Million)	FY 14 (Million)	FY 15 (Million)	FY 16 (Million)	Total per Component
Labor	\$42.90	\$44.65	\$45.87	\$47.45	\$49.60	\$51.45	\$53.22	\$335.14
Hardware	\$5.13	\$5.54	\$5.48	\$5.76	\$5.63	\$5.69	\$5.91	\$39.14
Software	\$0.00	\$0.00	\$0.00	\$3.35	\$5.87	\$5.99	\$6.12	\$21.33
Transport	\$7.60	\$18.39	\$18.85	\$19.31	\$19.73	\$20.14	\$20.60	\$124.62
TOTAL	\$55.63	\$68.58	\$70.20	\$75.87	\$80.83	\$83.27	\$85.85	\$520.23

For the thin-client solution, the cost components for both software and transport are similar to those for the thick-client solution. The reason for the similar software costs is because both solutions require a license per seat and the number of seats is the same. The same applies to the transport cost component, because the same number of seats are connected to the data centers and between the data centers, which make up the ONE-NET enterprise network. Similar to the thick-client solution, the cost of the first three years of software licenses is included in the initial purchase of the thin-clients as shown in Table 7. Even though the software resides on the servers instead of on the PCs, there is still a client access license (CAL) that applies to each seat.

Based on Table 7 values, the sum of the total thin-client costs for labor, hardware, software, and transport over the seven year period is \$520 million (rounded up).

Using the same VMware tool for calculating the desktop power and cooling power of the thick-client, the input for the thin-client is adjusted to the costs associated with the thin-client hardware. Table 8 provides the annual cost for power and cooling power for the thin-client. Table 8 also includes the power consumption required for the virtual desktop servers needed to support the thin-client solution, which is in addition to

what the current servers use to support the thick-client PCs. The proposed number of virtual servers to support the thin clients is 181 ESX servers.

Thus, the annual cost for operating and cooling power for the thin-client solution is shown in Table 8.

Table 8. Annual Cost of Power and Cooling on Thin-Client Solution
(From VMware, 2010)

Desktop Power and Cooling Costs	With VMware View (Projected) and Thin Clients		
	Virtual Desktop Server	Thin Client	Total
Hardware operating power (W/h)	750	15	765
Air conditioning cooling power (W/h)	600	12	612
Operating hours per day (h)	24	12	36
Total power consumption (W/d)	5,850,000	8,840,016	14,690,016
Total power and cooling (kW/yr)	2,106,000	3,182,406	5,288,406
Cost per kWh	\$0.1016	\$0.1016	\$0.106
Total annual cost	\$213,970	\$323,332	\$537,302

Based on a fixed annual cost of \$0.1016 per kilowatt hour (kW h) over the span of the contract, the total seven year cost for operating and cooling power for the current ONE-NET thin-client solution is \$3.8 million (rounded up).

The thin-client solution thus results in roughly \$11.6 million savings in the power and cooling cost, thereby offering a compelling advantage for its selection over the thick-client approach. This answers the third detailed question: Are there any other compelling advantages to deploying thin-client desktops and virtualization-based server technology over the current thick-client desktops?

The following additional cost analysis is performed to determine the up-front IT capital and operating investment to support the thin-client seats. The IT operating and capital investment costs include the costs for the virtualization servers, storage

infrastructure, and licenses for the VMware infrastructure. Table 9 shows the sum of total IT capital and operating investment from the provided VMware cost tables of APPENDIX C.

Table 9. IT Capital and Operating Investment (From VMware, 2010)

Expected Investment in Virtualization	Year 1 (M)	Year 2 (M)	Year 3 (M)	Year 4 (M)	Year 5 (M)	Total (M)
IT Capital Investment	\$8,012,287	\$9,627,480	\$11,504,809	\$2,709,958	\$2,786,107	\$34,640,641
IT Operating Investment	\$476,256	\$962,632	\$1,497,270	\$1,608,741	\$1,768,490	\$6,313,389
Total Investments	\$8,488,543	\$10,590,112	\$13,002,079	\$4,318,699	\$4,554,597	\$40,954,030

Given the seven-year TCO period, the proposed plan for a thin-client solution will take five years to complete. However, the current plan for rolling out thick clients is based on the same schedule. This is an important point, as it answers the fourth detailed question: Can the deployment of thin-client desktops be done during the current seven year technology support period and still lower the TCO?

3. TCO Comparison

Table 10 shows the complete TCO comparisons between thick-client and thin-client solutions. The second row in Table 10 includes the total costs for Labor, Hardware, Software, and Transport from the right column of Table 4. The Power & Cooling cost in the second row of Table 10 comes from Table 6. The third row of Table 10 includes the total costs for Labor, Hardware, Software, and Transport from the right column of Table 7. The Power & Cooling cost in the third row of Table 10 comes from Table 8, and the Virtualization cost is from Table 9. As shown in Table 10, the TCOs for the thick-client solution and the thin-client solution are, respectively, \$803 million and \$565 million.

Table 10. TCO Comparison between Between Thick-Client and Thin-Client Solutions

Alternative Solutions	Labor	Hardware	Software	Transport	Power & Cooling	Virtualization	TCO
Thick-Client	\$545	\$97	\$21	\$125	\$15.4	\$0	\$803
Thin-Client	\$335	\$39	\$21	\$125	\$3.8	\$41	\$565

The proposed thin-client solution thus results in a much lower TCO than does the current thick-client solution, by \$238 million over the seven-year period. As previously explained, the cost of software licenses and transport for the same number of seats is similar in both alternatives. While the thin-client solution does require \$41 million in additional investment for the virtualization resources in the data center to support the thin clients, the lower labor and hardware costs of the thin-client solution dominate the TCO differences. The deployment methodology is to use a similar deployment schedule for both the thin-client solution and the current thick-client solution. The methodology allows the additional ramp up in server resources to support the thin-clients at a lower TCO, which answers the fifth detailed question: What would be the recommended deployment methodology that will achieve the lower TCO? The incremental 30% deployment during the first three years and the last 10% over the following two years were used in the CPS and VMware calculation models for both thick-client and thin-client alternatives. Therefore, the lower TCO result for the thin-client solution supports the recommended proposal for the alternative solution.

D. CHAPTER SUMMARY

The Total Cost of Ownership (TCO) concept offers excellent insight into business processes and levels of service, which will help improve efficiencies and effectiveness. The object of good business is to attain the lowest TCO. In order to compare the TCO of the current thick-client solution with the proposed thin-client solution, separate calculations were done for the costs per seat over the seven-year period. For the thick-client solution, the CPS model was adjusted for inflation by 10.7% to cover the applicable period of FY 10 through FY 16. An additional VMware tool was used to

calculate the power and cooling costs for the thick-client solution, which were added to the CPS model results as the TCO for the thick-client solution.

For the proposed thin-client solution, a similar adjustment was done to the CPS model using the same 10.7% inflation increase, except that the inputs to the model for the hardware and labor, which were different from the thick-client inputs. The VMware tool was also used to calculate the power and cooling costs for the thin-client solution, which were added to the CPS model results, as well as the additional costs for the upfront IT capital and operating investment giving the TCO for the thin-client solution.

In calculations for both the thick-client and thin-client solutions, the same 27,284 seats were used. The potential savings of \$238 million is significant, and the proposed thin-client solution does not introduce much technical or schedule risk because the technology is already proven. Therefore, the answer is ‘yes’ to the second detailed question: Will the TCO of ONE-NET over the next seven years be low enough to justify the additional risk and up-front investment in deploying the thin-client desktops?

The main reasons for the cost savings are the much lower labor costs and the much lower power and cooling costs. The lower energy cost answers the third detailed question: Are there any other compelling advantages to deploying thin-client desktops and virtualization-based server technology over the current thick-client desktops? Also, the point of using the existing infrastructure and deploying the additional resources to support the thin clients makes it possible to use the same schedule as the current thick-client solution. Having the deployment effort for the thin-client solution priced by a vendor such as VMware within the same time frame scheduled for the thick-client solution provides the answer to the fourth detailed question of: Can the deployment of thin-client desktops be done during the current seven year technology support period and still lower the TCO? The TCO comparison between the thick-client and thin-client solutions is summarized in Table 10.

The same methodology for deploying thick clients is recommended for deploying thin clients, and since the TCO calculations using this recommendation show a lower TCO for the thin-client solution, the fifth detailed question, “What would be the

recommended deployment methodology that will achieve the lower TCO?” has been answered. The answer is a similar incremental 30% per year deployment during the first three years and the last 10% over the following two years.

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSION

Motivated by the prospects of reducing the TCO of ONE-NET, the hypothesis of this thesis is that by deploying a thin-client desktop and virtualization-based server solution, the TCO over the next seven years would be lower than that of the current plan for thick-client PCs. The primary question then is whether the TCO of the proposed alternative thin-client solution will be lower than that of the current thick-client solution. Answering the question amounts to answering the detailed questions. The detailed questions and their answers are summarized next.

A comprehensive review of thin-client technology indicates that the thin-client and server virtualization technology has finally matured and solved the issues of GUI responsiveness similar to the desktop PCs, and is therefore technically feasible as an alternative to the current thick-client solution. The answer to the first detailed question, “Will replacement of the current thick-client desktops with thin-client desktops meet the operational requirements of ONE-NET?” is that the same software image building process is used for the thin-client as for the thick-client, so the operational requirements for applications and security are the same. In fact, the security of the thin-client solution is much more robust than that of the thick-client solution, because the former is all done centrally in the data center and is not exposed to various implementations and enforcement at the end user seats.

The second detailed question is: Will the TCO of ONE-NET over the next seven years be low enough to justify the additional risk and up-front investment in deploying the thin-client desktops? The calculations indicate that the additional investment of \$41 million is substantially less than \$210 million in labor cost savings and an additional \$58 million in hardware savings of deploying thin-client seats, instead of the refreshed thick-client PCs.

An interesting aspect of the thin-client solution is the advantage of lower operating and cooling power costs compared to those of the thick-client solution. Based on the findings in this thesis, the power and cooling cost for the thin-client solution is

\$11.6 million less than that of the thick-client solution. Thus the thin-client solution has a compelling advantage over the thick-client desktop PCs from an energy savings perspective. This compelling advantage answers the third detailed question: Are there any other compelling advantages to deploying thin-client desktops and virtualization-based server technology over the current thick-client desktops?

The recommendation to use the existing infrastructure and deploying the additional resources to support the thin clients is supported by a lower TCO, as compared to that of the thick-client solution (Table 10). It is thus possible to use the same schedule used for the current thick-client solution, thereby answering the fourth detailed question of this thesis: Can the deployment of Thin-Client desktops be done during the current seven year technology support period and still lower the TCO? The proposed thin-client solution is feasible within the given timeframe. Furthermore, the recommendation for deploying the thin-client solution using the same incremental 30% per year deployment during the first three years and the last 10% over the following two years answers the fifth detailed question: What would be the recommended deployment methodology that will achieve the lower TCO? The resulting lower TCO supports this recommendation and provides a strong incentive to pursue the thin-client solution.

The latest technology of cloud computing offers an interesting topic for future work in studying how the concept relates to enterprise networks and TCO. The main paradigm of cloud computing is that the thin-client user does not know which physical server in the network ‘cloud’ it is communicating with. In other words, the traditional server is replaced with an abstract service that can be running on several different physical servers that synchronize their data.

While many institutions, including the U.S. Navy, are looking for ways to take advantage of cloud computing technology, there are several serious issues that still need to be resolved such as IA. All Navy IA Controls and certification processes today rely heavily on knowing the exact network topology, installed software location, and end-to-end communication channels. Resolving the IA challenges of cloud computing is also a potential topic for future research.

APPENDIX A. CPS MODEL CALCULATION RESULTS FOR FY08 UNADJUSTED THICK-CLIENT SOLUTION

Table 11 provides an excerpt from the CPS model calculations for the three regions: the island nation of Bahrain, Far East, and Europe.

Table 11. FY 08 Unadjusted Thick-Client Solution Cost Results

2008								
Summary Cost Per Seat Data	Seats	Total Cost	LNSC Labor Cost	TNOSC Labor Allocation	Other Labor Cost	Hardware Cost	Software Cost	Transport Cost
Bahrain CPS 2008	2,205	\$13,350,590	\$7,271,791	50	\$293,754	\$1,009,097	\$244,811	\$4,531,136

Yokosuka 2008	7,707	\$19,178,472	\$1,894,884	\$10,045,198	\$98,661	\$3,483,683	\$855,671	\$2,800,374
Guam 2008	2,838	\$8,801,376	\$3,438,984	\$3,700,314	\$63,151	\$1,283,727	\$315,200	50
Atsugi 2008	3,555	\$9,438,701	\$1,753,526	\$4,633,538	\$58,840	\$1,657,250	\$394,695	\$940,851
Sasebo 2008	733	\$4,750,751	\$1,188,586	\$955,382	\$28,573	\$331,715	\$81,381	\$2,165,114
Misawa 2008	762	\$2,768,493	\$1,230,353	\$993,180	\$31,355	\$338,845	\$84,601	\$90,158
Okinawa 2008	787	\$2,393,314	\$1,464,023	\$1,025,765	\$54,162	\$361,987	\$87,377	50
Diego Garcia 2008	241	\$1,417,609	\$948,941	\$314,116	\$21,969	\$105,826	\$26,757	50
Singapore 2008	374	\$1,842,708	\$965,752	\$487,466	\$22,148	\$196,283	\$41,523	\$129,534
Korea 2008	925	\$4,018,042	\$1,361,200	\$1,205,632	\$37,548	\$435,284	\$102,698	\$875,679
Far East CPS 2008	17,923	\$ 56,209,465	\$ 14,246,260	\$ 23,360,693	\$4,16,407	\$ 8,194,601	\$ 1,989,904	\$7,001,710

Naples 2008	3,072	\$13,856,319	\$2,212,134	\$6,003,047	\$114,720	\$1,412,585	\$341,069	\$3,772,763
Sigonella 2008	1,913	\$8,006,234	\$2,546,586	\$3,738,226	\$147,799	\$826,392	\$212,391	\$534,841
Rota 2008	1,219	\$7,139,098	\$2,853,161	\$2,382,069	\$142,104	\$565,835	\$135,340	\$1,060,590
London 2008	-							
La Maddalena 2008	370	\$2,436,730	\$1,158,782	\$723,023	\$60,108	\$197,897	\$41,079	\$255,840
Souda Bay 2008	582	\$3,168,755	\$1,331,736	\$1,137,296	\$58,026	\$294,764	\$64,617	\$282,316
Europe CPS 2008	7,156	\$ 34,807,136	\$ 10,102,398	\$ 13,983,661	\$522,758	\$ 3,297,472	\$ 794,496	\$5,906,350

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX B. REGIONAL ROLLUP CALCULATIONS FOR TCO OF THICK-CLIENT SOLUTION

Table 12 provides the rollup calculation results from the thick-client data output of the costs per seat model. For each year, the individual cost components for labor, hardware, software, and transport were summed up by the three regions of the island nation of Bahrain, Far East (FE), and Europe (EU).

Table 12. FY 05 Through FY 11 Regional Rollup Costs

FY 05	Labor	Hardware	Software	Transport
Bahrain	\$6,843,128.98	\$949,033	\$0	\$2,890,494
FE	\$34,353,139.14	\$7,694,836.58	0	\$2,877,569
EU	\$22,258,738.40	\$3,005,197.91	0	\$1,094,918.00
total	\$63,455,006.52	\$11,649,067.72	\$0	\$6,862,981.00

FY 06	Labor	Hardware	Software	Transport
Bahrain	\$7,075,801	\$968,014	\$0	\$4,346,660
FE	\$35,541,323.17	\$7,998,673.32	0	\$6,703,797
EU	\$23,015,563.17	\$3,191,781.86	0	\$5,560,265.34
total	\$65,632,687.25	\$12,158,469.08	\$0	\$16,610,722.26

FY 07	Labor	Hardware	Software	Transport
Bahrain	\$7,316,573	\$988,342	\$0	\$4,437,940
FE	\$36,751,442	\$8,036,467.96	0	6,857,698.67
EU	\$23,798,845	\$3,229,649.52	0	5,729,455.99
total	\$67,866,860	\$12,254,459.67	\$0	\$17,025,094.40

FY 08	Labor	Hardware	Software	Transport
Bahrain	\$7,565,545	\$1,009,097	\$244,811	\$4,531,136
FE	\$38,023,250	\$8,194,600.88	\$1,989,904.28	\$7,001,710
EU	\$24,608,817	\$3,297,472.16	\$794,496.18	\$5,906,350
total	\$70,197,612	\$12,501,170.42	\$3,029,211	\$17,439,196.91

FY 09	Labor	Hardware	Software	Transport
Bahrain	\$7,822,817	\$1,030,288	\$428,488	\$4,626,290
FE	\$39,338,538.86	\$8,392,742.36	\$3,482,901.03	\$7,162,425
EU	\$25,445,675.33	\$3,275,527.09	\$1,390,595.31	\$6,030,383

total	\$72,607,031.07	\$12,698,557.87	\$5,301,985	\$17,819,098.84
--------------	------------------------	------------------------	--------------------	------------------------

FY 10	Labor	Hardware	Software	Transport
Bahrain	\$8,088,927	\$1,051,924	\$437,487	\$4,723,442
FE	\$40,699,560	\$8,529,087	\$3,556,041.95	\$7,312,836
EU	\$26,311,356	\$3,331,012.16	\$1,419,797.81	\$6,157,021
total	\$75,099,843	\$12,912,023.57	\$5,413,326	\$18,193,299.92

FY 11	Labor	Hardware	Software	Transport
Bahrain	\$8,363,991	\$1,074,015	\$446,674	\$4,822,635
FE	\$42,107,392	\$8,721,778.09	\$3,630,718.83	\$7,473,535
EU	\$27,206,092.37	\$3,485,840.44	\$1,449,613.57	\$6,306,689.44
total	\$77,677,475.77	\$13,281,633.42	\$5,527,006	\$18,602,859.37

APPENDIX C. VMWARE CAPITAL AND OPERATING INVESTMENT COSTS FOR VIRTUALIZATION

Table 13 provides the calculated costs from VMware Inc. based on the 27,284 seats required by ONE-NET.

Table 13. VMware Virtualization Capital and Operating Investment Costs

VMware TCO/ROI Calculator Report						
View	Initial / Year 1	Year 2	Year 3	Year 4	Year 5	Total
IT Capital Investment						
VMware View Server and Storage Infrastructure Configuration	\$2,307,555	\$2,307,555	\$2,999,822	\$761,494	\$837,643	\$9,214,069
VMware View Configuration - Desktop Virtualization	\$2,545,732	\$3,845,025	\$4,714,187	\$1,000,764	\$1,000,764	\$13,106,472
VMware View Configuration - Thin Client Migration	\$3,159,000	\$3,474,900	\$3,790,800	\$947,700	\$947,700	\$12,320,100
Total IT Capital Investment	\$8,012,287	\$9,627,480	\$11,504,809	\$2,709,958	\$2,786,107	\$34,640,641
IT Operating Investment						
VMware View Support and Subscription	\$438,750	\$921,375	\$1,452,263	\$1,597,489	\$1,757,238	\$6,167,115
VMware View Implementation Professional Services and Labor	\$37,506	\$41,257	\$45,007	\$11,252	\$11,252	\$146,274
Total IT Operating Investment	\$476,256	\$962,632	\$1,497,270	\$1,608,741	\$1,768,490	\$6,313,389
Total Investments	\$8,488,543	\$10,590,112	\$13,002,079	\$4,318,699	\$4,554,597	\$40,954,030

© Copyright 2001 - 2009, VMware, Inc. and Alinean, Inc. All rights reserved.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Bach, M. J. (1986). *Design of the UNIX operating system*. New York: Prentice Hall.
- CSC. (2009, April 23). Management plan/technical approach, I-D-30-55.
- Esposito, R., & Slack, K. (2010). Three generation of thin-clients. ProSync Web site, Retrieved on December 8, 2009, from
<http://www.prosync.com/PDF/thinclientevolution.pdf>
- Graham, B. (2009, July 23). Tivoli scripting for automatic scan and patch for PC refresh systems, SPAWAR Systems Center-Pacific, San Diego, 1–2.
- Hanada, M. (2009, June 9). System security architecture description (SSAD) staging VLAN COI, VLAN, SPAWAR Systems Center-Pacific, San Diego, 4–5.
- Lelfwitch, J. (2009, February 2). Build process for WBSC 118 image ONE-NET, SPAWAR Systems Center-Pacific, San Diego, 9–11.
- Martinez, J. (2010, April 20). Altiris enterprise deployment solution technical design, 13–14.
- McDaniel, D., & Falcone, N. (2008, May 2). System security architecture descriptions (SSAD) transport community of interest (COI), 4–5.
- McKuick, M. (1999, January). Open sources: voices from the open source revolution. Wikipedia Web site, Retrieved on July 20, 2010 from
http://en.wikipedia.org/wiki/Berkeley_Software_Distribution
- Mui, L., & Pearce, E. (1992). The definitive guide to the X window system, volume eight, X window system administrator's guide for X11 release 4 and release 5, O'Reilly & Associates, Inc.
- NNWC. (2001, February 23). ONE-NET IT service delivery description and service levels (SLA), 10–11.
- NNWC. (2005, December 20). ONE-NET cost per seat (CPS).
- PC Magazine. (2010). Encyclopedia definition: enterprise network. PCMAG Web site, Retrieved on July 20, 2010, from
http://www.pcmag.com/encyclopedia_term/0,2542,t=enterprise+network&i=42648.00.asp
- Shimoko, S. (2008a, May 1). WPMS automated inventory, 2.

- Shimoko, S. (2008b, May 1). Tivoli windows patch management system (WPMS) enterprise support procedures, 8–9.
- Solution Matrix. (2010). Total cost of ownership analysis (TCO)/cost of ownership (COO). Solutionmatrix Web site, Retrieved on July 20, 2010, from <http://www.solutionmatrix.com/total-cost-of-ownership.html>
- SSC PAC. (2009a, February 4). Workstation baseline, WBSC ver 118.
- SSC PAC. (2009b, February 19). ONE-NET windows patch management process (WPMP).
- SSC PAC. (2010a, March 13). OCONUS navy enterprise network (ONE-NET) overview, 8.
- SSC PAC. (2010b, July 29). Project review—PC refresh, project brief.
- SSC PAC. (2010c, July 29). Server farm refresh plan of action and milestones (POAM).
- Tech Terms. (2010). Thick client. Techterms Web site, Retrieved on July 20, 2010 from, <http://www.techterms.com/definition/thickclient>
- VMware. (2010). VMware view 4 built for desktop. VMware Web site, Retrieved on May 10, 2010, from <http://www.vmware.com/support/pubs>

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California
3. Dr. Thomas V. Huynh
Naval Postgraduate School
Monterey, California
4. Ms. Sandra Baekey
Naval Network Warfare Command
Norfolk, Virginia
5. Mr. Joseph Gildersleeve
Naval Network Warfare Command
Norfolk, Virginia
6. Mr. Thomas D. Scruggs
Program Executive Office—Enterprise Information Systems
Arlington, Virginia
7. Mr. Jack Cabana
SPAWARSCEN PACIFIC
San Diego, California